



EUROPEAN PATENT APPLICATION

(43) Date of publication:
07.08.2002 Bulletin 2002/32

(51) Int Cl.7: H04L 12/24

(21) Application number: 02394018.2

(22) Date of filing: 05.02.2002

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR
Designated Extension States:
AL LT LV MK RO SI

(72) Inventors:
• Lam, Trinh T.
Santa Clara, CA 95050 (US)
• Turner, Doug G.
Gloucester, Ontario K1B 4A9 (CA)

(30) Priority: 06.02.2001 US 776852

(74) Representative: Boyce, Conor et al
F. R. Kelly & Co.,
27 Clyde Road,
Ballsbridge
Dublin 4 (IE)

(71) Applicant: Nortel Networks Limited
St. Laurent, Quebec H4S 2A9 (CA)

(54) Service level agreement manager for a data network

(57) An SLA manager suitable for assessing the performance of a data network having a transport layer, a services layer and an applications layer. The SLA manager has an input to receive performance metrics data from probes in the data network. The input also re-

ceives an SLA profile that defines one or more performance objectives related to the data network that the SLA manager is to track. The SLA manager measures the performance objectives defined by the SLA profile on the basis of performance metrics generated from probes associated with different layers of the data network.




FIG. 1

Description

[0001] The invention relates to devices and methods useful in the assessment of the level of service delivered to a customer against performance objectives relating to the operation of a data network.

[0002] Communications networks, especially packetised data networks, are currently utilised in various applications for transmission and reception of data between parties at different locations. A typical data transmission system includes a plurality of end-user sites and a data network, which resides between the sites to facilitate communications.

[0003] Typically, when a party needs to send and receive data over distances, the party (end-user) enters into a service contract with a service provider to provide access to a data network. Depending on an individual end-user's needs, the service contract may include provisions that guarantee a certain minimum performance requirements that the service provider must meet. For example, if the end-user expects to send and receive a certain amount of data on a regular basis, the end-user may want the service provider to guarantee that a certain minimum bandwidth will be available to the end-user at all times. Certain end-user applications are sensitive to transmission delays and/or the loss of data within the network. Specifically, while loss of data packets can generally be detected by end-users, and lost packets can be retransmitted, certain applications cannot function when the percentage of lost data exceeds a given level. Thus, the end-user may want the service provider to guarantee that the average or minimum ratio of data units delivered by the network to data units offered to the network at the far-end is above a certain percentage and/or that the average or maximum transmission delays will not exceed a certain duration.

[0004] From a service provider's perspective, it would be competitively advantageous to be able to demonstrate to potential and existing end-users that the service provider is capable of meeting and does meet such performance objectives. Thus, the capability to provide analysis of the level of service delivered by the end-user would be advantageous from both an end-user and service provider's standpoint.

[0005] Various devices have been proposed in the past to provide some measure of the performance of a data network. The existing devices only track the performance of some limited aspects of the data network, which may not be sufficient to assess the overall network performance. Accordingly, there is a need in the industry to develop a system and components of a system that can assess the performance of a data network on a broader basis such as to provide a more comprehensive image of the data network performance.

[0006] In a broad aspect, the invention provides a computer readable storage medium containing a program element suitable for execution by a computing entity for implementing a service level agreement (SLA) manager. The SLA manager is suitable for assessing the performance of a data network having a transport layer, a services layer and an applications layer. The data network also has a plurality of probes associated with different layers of the data network and generating performance metrics data.

[0007] The SLA manager receives the performance metrics data and an SLA profile associated to a user of the data network. The SLA profile defines one or more performance objectives that the SLA manager is to track. The SLA manager measures at least one of the performance objectives defined by the SLA profile at least in part on the basis of performance metrics data from a first probe associated with one layer of the data network in conjunction with performance metrics data from a second probe associated with another layer of the data network.

[0008] The advantage of the SLA manager resides in its ability to track complex performance objectives that can be assessed on the basis of performance metrics issued by probes associated with different network layers. The resulting view of the data network performance is more complete than the results obtained from a device designed to track performance objectives on the basis of performance metrics issued by a probe(s) associated to the same network layer.

[0009] The transport, services and applications layers reflect a hierarchical organization of the data network. The transport layer is the lowest layer in the hierarchy, the applications layer the highest layer in the hierarchy and the services layer is the intermediate layer in the hierarchy. Typically, the transport layer refers to the functional elements of the data network that are responsible for the transport of data from one point of the network to another point of the network. Optionally, the transport layer can be divided in two sub-layers, namely the transport core layer and the virtual transport layer. The transport core layer typically encompasses the hardware (physical level connections) and the protocols that govern the hardware connections between network nodes. Examples of such protocols include the Internet Protocol (IP), the Asynchronous Transfer Mode (ATM), the Multi-Protocol Label Switching (MPLS), the Frame Relay, and the Time Division Multiplexing (TDM), and emerging native optical services such as managed Lambda services. The virtual transport layer is a one-step abstraction toward the services layer. The virtual transport layer encompasses software entities that act as interfaces between the core transport layer and the services layer and the applications layer. A virtual transport layer employs the underlying services of the transport core in delivery of service. For example, a virtual transport layer can be an enterprise IP VPN comprised of virtual IP routers running over an ATM infrastructure. Alternatively, a virtual transport layer can provide a dial access service with user authentication, address assignment, and secure tunneling to an enterprise site.

[0010] The services layer encompasses facilities that form a basis upon which specific user oriented applications can run. Typically, the services layer includes facilities that are most of the time transparent to the user but focused on

delivery of services for the user. Examples include security services such as user authentication and public-key based encryption, subscriber-based firewalls, policy services (controlling aspects of the network operation based on a certain policy again oriented to the user/subscriber), billing services, IP-customized services based upon aspects such as time of day, service activation based upon subscriber profiles and class of service, customer self-service functions such as dynamic bandwidth request services, etc.

[0011] Finally, the applications layer encompasses the functions of the network with which the user interacts directly. Examples include software entities that the user will specifically invoke to perform a certain task, such as spreadsheet applications, voice mail applications, multimedia applications, wireless telephony applications; a range of collaborative applications including gaming, remote learning, tele-medicine, enterprise video meet services, remote security and monitoring services, etc.

[0012] Other broad aspects of the invention include a method to assess the performance of a data network and a tool to generate an SLA profile.

[0013] A detailed description of examples of implementation of the present invention is provided hereinbelow with reference to the annexed drawings, in which:

Fig. 1 is a functional block diagram of a tool for creating an SLA profile, in accordance with a specific example of implementation of the present invention;

Fig. 2 is a more detailed block diagram of an SLA performance metrics selector unit and an SLA function selector unit of the tool shown in Fig. 1;

Fig. 3a is a more detailed block diagram of an SLA rule selector of the tool shown in Fig. 1;

Figure 3b is a flowchart that illustrates the process of the tool shown in Fig. 1;

Fig. 4 is a functional block diagram of an SLA manager in accordance with an example of implementation of the present invention;

Fig. 5 is a high-level functional block diagram of an action thread of the SLA manager shown in Fig. 4;

Fig. 6 is a variant of the action thread shown in Fig. 5;

Fig. 7 is a block diagram of another variant of the action thread shown in Fig. 5;

Fig. 8 is a block diagram of two action threads linked together;

Fig. 9A is a further variant of the action thread shown in Fig. 5;

Fig. 9B is a further variant of the action thread shown in Figure 5;

Fig. 10 is a block diagram of a complete system using the SLA manager to measure the performance of a data network;

Fig. 11 is a block diagram of a data network having a transport functional layer, a services functional layer and an applications functional layer;

Fig. 12 is a flowchart illustrating the operation of the SLA manager; and

Fig. 13 is a structural block diagram of a computing device for executing the software that implements the tool for creating an SLA profile or the SLA manager.

[0014] Figure 11 provides an illustration of a typical data network. The data network 100 provides various services to seven users (arbitrary number) namely user 112, 114, 116, 118, 120, 122 and 124. The users run different applications. User 112 runs a web browsing software. User 114 runs a graphics software. User 116 runs a video capture software. User 118 is recording a voice mail message at a wireless telephone. User 120 is engaged in a conversation with a remote party at a wireless telephone. User 122 uses a word processing application. User 124 runs a game application.

[0015] The data network 100 defines three functional layers namely a transport layer 126, a services layer 128 and

an applications layer 130. The transport layer 126 is divided in two parts, namely a core transport layer 132 and a virtual transport layer 134. Further layering of the network into multiple hierarchical transport domains is possible: this representation is not intended to be limiting.

[0016] The data network 100 uses an SLA manager (not shown in Figure 11) to track the performance of the data network 100 based on performance objectives specified in an SLA profile. The operation of the SLA manager will be described in greater detail later.

[0017] The SLA profile directs the operation of the SLA manager by defining the performance objectives to be tracked. Different users 112-124 are likely to require different SLA profiles to meet individual needs.

[0018] Fig. 1 is a high-level functional block diagram of a tool 10 for creating an SLA profile. The tool 10 is implemented in software on a suitable computing platform that will be described later. The tool 10 receives user commands and displays information to the user through a user interface 12. A plurality of functional units are provided allowing the user to specify through the user interface 12 the desired characteristics or features of the SLA profile that is to be created. In the example of implementation shown in Fig. 1, the tool 10 has three such functional units. This is not an essential requirement of the invention as some of these units can be omitted, or additional units can be added without departing from the spirit of the invention.

[0019] The three functional units of the tool 10 are the SLA performance metrics selector 14, the SLA function selector 16 and the SLA rule selector 18. The SLA performance metrics selector 14 allows a user to specify through the user interface 12 SLA performance metrics. A performance metrics is a measurable factor that defines one or more aspects of the data network operation. Such factor will normally have a range of variations.

[0020] Fig. 2 illustrates in greater detail the SLA performance metrics selector 14. Under one possibility, the user is presented with a menu listing a plurality of predetermined SLA performance metrics and the user can select in the list a particular performance metrics of interest. Alternatively, the user can enter in a field the name of a certain performance metrics of interest instead of selecting it through a menu system.

[0021] Probes in the data network 100 generate performance metrics data. The performance metrics data is the measurement of a performance metrics. For example, the performance metrics data represents the value of a performance metrics (the measurable parameter) at a given time. Consider the more specific example of the bandwidth over a certain network segment. Here the performance metric is the bandwidth and the performance metrics data is the value of the bandwidth at a certain time, location, etc.

[0022] A probe for generating performance metrics data may use a combination of technologies. For example, a probe may use FTP to retrieve a remote file stored as ASCII text using XML for structuring the contents. Alternatively, a similar ASCII file could be transferred via FTP but be formatted using a CSV (Comma Separated Value) format. Some common probes can be created to manage standard protocols, such as SNMP, and others need to be created for proprietary data schemes. Note also that more than one type of probe may be used for collecting data for a given service. The data required may be distributed across many sources within a network (or other system), and require different probes for extracting the data.

[0023] Referring back to Figure 11, the data network 100 has a plurality of probes that deliver performance metrics data that are associated with the various functional layers of the data network 100. More specifically, the data network 100 includes a set of probes associated with the transport layer 126, a set of probes associated with the services layer 128 and a set of probes associated with the applications layer 130. In general, a probe associated with a certain functional layer of the data network 100 will deliver performance metrics data relating to the operation of the that functional layer.

[0024] In Figure 11, the probes are shown as circles. A circle with the letter "T" identifies a transport layer probe, the number that follows the letter "T" represents the probe number. Similar representations are used for the services layer and the applications layer probes with the exception that the probes are identified by the letters "S" and "A", respectively.

[0025] The SLA performance metrics selector 14 can accept as input from the user the identification of the SLA performance metrics of interest for the purpose of a particular SLA profile. This approach requires a mapping between the SLA parameters and the corresponding probes, such that in use data only from the probes producing the performance metrics data of interest will be considered. The mapping can be accomplished through a data structure that associates the SLA performance metrics to the respective probes in the data network 100.

[0026] Alternatively, the SLA performance metrics selector 14 can accept as input the identification of the probes in the network that generate the SLA performance metrics data of interest. This method indirectly specifies the SLA performance metrics required for a given profile.

[0027] The SLA function selector 16 allows the user to define a function to be applied on performance metrics data. The SLA function can be any suitable function applied on performance metrics data. Examples include:

- a) minimum function - this represents a function, which determines the smallest value in a set of performance metrics data;
- b) maximum function - this represents a function, which determines the largest value in a set of performance

metrics data;

c) average (mean) - this represents a function that adds all the values in a set of performance metrics data and then divides it by the number of values in the performance metrics data;

d) summation - this represents a function that adds all the values in a set of performance metrics data;

e) maximum consecutive - this represents the highest number of consecutive samples in a set of performance metrics data with a value of "True". This function can be used to evaluate clauses such as "more than two days in a row";

f) other Boolean functions; and

g) heuristics functions.

[0028] The heuristics functions are useful in refining the data stream that is generated by the probes such as to limit the amount of data to be analysed and retain only the information that is the most useful in the context of a given SLA profile. As discussed earlier, the SLA performance metrics selector 14 determines or defines a set of probes in the data network 100 applicable to a given SLA profile. These probes define a set of paths through the data network 100 over which measurements are to be taken during the execution of the SLA profile. Conceptually, if a tree is drawn to identify all the paths possibilities, a very large set of probe paths will be produced, particularly in the case where the number of probes selected is large. The heuristics functions allow the user to specify criteria to be applied during the processing of the performance metrics data generated by the probes when the SLA profile is run. Such criteria would normally have the effect of reducing the amount of data to be processed, in effect "pruning" the full tree of probe paths. A simple example of a criterion is to use a measurement from one probe one day and use the measurement from another probe another day. The criteria may also involve fully blocking the information from one or more probes, if that information is of no interest to the particular SLA profile.

[0029] Consider now a more complex example. In Figure 11, it is possible to derive a tree denoting a set of probe paths. Based upon SLA requirements, only subset of paths and sub-paths through the tree are of interest for a specific SLA profile. For example, an SLA performance metrics based on path A1-S1-S4-T1-Tn-S6-A7 can be employed to measure the performance of the data network from the perspective of the user 112. Alternatively, a sub-path to abstract the details of the transport layer 126 can be employed. Such "reduced" SLA performance metrics could be based on path A1-S1-S4-S6-A7.

[0030] An SLA performance metrics data processed by an SLA function produces a composite SLA data. This is illustrated in Fig. 2. Note that an element of information processed by an SLA rule (the SLA rules will be described later) can also produce a composite SLA data. The SLA performance metrics selector 14 allows the user to select one or more SLA performance metrics data, as described above. The SLA function selector 16 allows the user to specify a function to be applied on the performance metrics data generated by one or more probes during the execution of the SLA profile by the SLA manager.

[0031] The choice of the SLA function can be made through a menu system, typing in a field the name of the SLA function, or by any other suitable selection mechanism. As shown by the dotted lines in Fig. 2, the SLA performance metrics #1 can be linked to any one of the SLA functions. The same functional linkage can also be established between any one of the remaining SLA performance metrics in the list. This has not been shown in the drawings for simplicity. It should be appreciated that it is not essential that every SLA function be applicable to every SLA performance metrics. Possibilities where an SLA function is applicable to only a subset of the entire set of SLA performance metrics are clearly within the scope of the present invention. In this case, the tool 10 includes a mechanism to verify the validity of the selection made by the user and avoid instances where the user selects an SLA function that cannot support a previously selected SLA performance metrics. It is within the reach of a person skilled in the art to create such validation mechanism and no further details are required here.

[0032] A composite SLA data is produced when an SLA function is applied on performance metrics data. In the example shown in Fig. 2, two such composite SLA data are produced and associated with SLA functions #1 and #2. It will be plain to a person skilled in the art that many other possibilities exist without departing from the spirit of the invention.

[0033] Fig. 3A provides a more detailed illustration of the SLA rule selector 18. The SLA rule selector 18 allows the user to specify one or more SLA rules to be applied on performance metrics data or to be applied on composite SLA data. An SLA rule when applied on data triggers at least one SLA action. An SLA action is either a positive action (some event takes place) or a null action (nothing happens). Note that the above does not imply that the application of an SLA rule produces an SLA positive action or a null SLA action every time. Possibilities where the application of an SLA rule produces every time an SLA positive action are within the scope of this invention. For example, when a condition specified by an SLA rule is met, SLA positive action #1 is triggered, otherwise SLA positive action #2 is triggered.

[0034] In one specific example, SLA rules can be Boolean functions. Those include:

- a) less than - the result is "True" if the data on which the SLA rule is applied is determined to be less than a certain mark, "False" otherwise;
- b) greater than - the result is "True" if the data on which the SLA rule is applied is determined to be greater than a certain mark, "False" otherwise;
- 5 c) less than or equal to - the result is "True" if the data on which the SLA rule is applied is determined to be less than or equivalent to a certain mark, "False" otherwise;
- d) equal to - the result is "True" if the value on which the SLA rule is applied is determined to be equivalent to a certain mark, "False" otherwise; and
- 10 e) between upper and lower boundaries - this represents a function that determines if data is between an upper boundary and a lower boundary.

[0035] The above list is not exhaustive and many other possibilities exist.

[0036] The flowchart at Fig. 3B summarises the process for the creation of the SLA profile by using the tool 10. At step 150, the process starts. At step 102, the user selects through the user interface 12 one or more SLA performance metrics. At step 104, the user selects one or more SLA functions and at step 106, the user selects one or more SLA rules. At step 108, the tool 10 generates through an output 20 (shown in Fig. 1) the SLA profile 22. The SLA profile 22 is data recorded on a computer readable storage medium directing the operation of an SLA manager 24.

[0037] When the SLA manager 24 runs, it implements a process referred to as "action thread". An action thread includes the processing of performance metrics data according to the SLA function(s) and/or SLA rule(s) defined by the SLA profile to produce SLA action(s). Note that an action thread may include additional steps without departing from the spirit of the invention. In the example shown in Fig. 4, the SLA manager 24 implements three action threads 26, 28, 30, it being understood that more or less action threads are permitted without departing from the spirit of the invention.

[0038] An advantageous feature is the ability to build with the tool 10 SLA profiles that define action threads measuring performance metrics based on information generated from probes associated to different functional layers of the data network 100. For instance, such an action thread can process performance metrics data from one or more probes associated with the transport layer 126 in conjunction with performance metrics data from one or more probes associated with the services layer 128. Evidently, other permutations are possible without departing from the spirit of the invention.

[0039] A detailed example of an action thread is shown in Fig. 5. The action thread includes an SLA rule that is applied to performance metrics data to generate an SLA action. There is no SLA function involved. In this case, the SLA rule compares the performance metrics data to the performance objective and, if the result of the comparison indicates that the performance objective has not been met, then an SLA action is triggered. Under this example, the SLA action is the generation of a message to signal a rebate event. That message can be directed to the electronic account of the customer to make the rebate effective.

[0040] Fig. 6 is a more complex example of an action thread where the performance metrics data is first subjected to an SLA function that generates composite SLA data to which is then applied an SLA rule.

[0041] Fig. 7 illustrates yet another example of an action thread characterised by the consecutive application of two SLA functions to performance metrics data. The resulting composite SLA data is then processed by an SLA rule.

40 [0042] Fig. 8 shows an example where two action threads share a common SLA function. Here, the action thread 30 includes an SLA function 32 that generates composite SLA data supplied to an SLA rule 34. The SLA rule 34 then generates an SLA action #1. The action thread 36 also comprises the SLA function 32 and further includes the SLA function 38 and the SLA rule 40 that generates a SLA action #2.

45 [0043] Fig. 9A is another example of an action thread where the SLA function includes a looping mechanism to return the result of the processing of the SLA function to its input so that it is processed again a number of times.

[0044] Fig. 9B is yet another example of an action thread. Here the application of an SLA rule produces an SLA action and also composite SLA data that could be used for further processing. In other words, the SLA rule acquires here the behaviour of an SLA function in addition to its basic functionality of SLA action generation. The composite SLA data generated by the SLA rule is suitable for processing by an SLA function or by the same or different SLA rule.

50 [0045] Fig. 10 is a block diagram of the SLA manager 24 that collects information from a data network. The SLA manager 24 connects to the probes T1-Tn, S1-Sn and A1-An that generate performance metrics data supplied to input 52 of the SLA manager 24.

[0046] The SLA manager 24 also receives through input 52 data representative of SLA profiles namely the SLA profiles 54, 56 and 58. Each profile may be associated to a different client or user of the data network 100. The SLA profiles 54, 56 and 58 may reside in a database that is remote or local to the SLA manager 24.

55 [0047] The SLA manager 24 includes a processing service 61 that processes the information received at the input 52. In order to determine if a certain performance objective is met, the SLA manager 24 will consider the performance metrics data generated by the probes T1-n, S1-n and A1-n. as described earlier, this determination is made by imple-

menting one or more action threads. The action threads give rise to SLA actions. In the example illustrated in Fig. 10, one SLA action is a positive SLA action which is the printing of a report on a printer 55, another SLA action is a positive SLA action which is the creation, modification or deletion of a billing record in a database 57 and another SLA action is also a positive SLA action which is a command to an e-mail server 59 to send an e-mail message to a recipient. A positive SLA action can be such as to trigger pro-active behaviour designed to alter the service provided to the end-user to automatically adjust the service to meet SLA performance objectives.

[0048] While assessing if a certain performance objective is met, the SLA manager 24 may consider the performance metric data generated by a single probe, by a plurality of probes belonging or associated to the same functional layer of the data network 100, or by a plurality of probes belonging to different functional layers of the data network 100.

[0049] The following table provides detailed examples of performance metrics that can be tracked in a data network. This example should be treated as an aid to understanding the concept of the invention and not as limiting in any way to the inventive principle. The person skilled in the art will readily appreciate that additional performance objectives can be added to this table or performance objectives in the table can be removed without departing from the spirit of the invention.

[0050] In the table:

1. The "class" column identifies to which class the performance metrics belongs. This categorisation is done to facilitate the analysis of the information and it is not essential in any way to the invention. Other systematic grouping can be considered without departing from the spirit of the invention
2. The "sub-class" column identifies a sub-class where the performance metrics belongs;
3. The "applicable probes" column identifies from which probes performance metrics data is obtained. Here, "T", "S" and "A" refer to transport layer probes, service layer probes and application layer probes, respectively;
4. When the column "service provider" is checked, the information related to the performance objective is useful primarily to the service provider; and
5. When the column "user" is checked, the information related to the performance objective is useful primarily to the user.

Class	Sub-class	Performance Metrics	Applicable Probes	Specific Example	Service provider	User
Availability	Transport	Network Element Mean Time Between Failure (MTBF)	S, T	% router uptime = 99.999%	✓	
Availability	Transport	Segment MTBF (e.g., ATM VC, Data Link Connection Identifier (DLCI), TDM channel)	S, T	Denver-Chicago connection uptime = 99.999%	✓	
Availability	Transport	Mean Time to Repair (MTTR) (for Network Elements, segments)	S, T	Downed trunks will be repaired within 2 hrs	✓	
Availability	Transport	Unscheduled Network Element Down Time	S, T	Max possible unscheduled router downtime per year = 4.5 hrs	✓	
Availability	Transport	Unscheduled Network Connections Down Time	A, S, T	Max possible unscheduled RNC-Quest downtime per year = 30 min	✓	
Availability	Transport	% IP destination address reachable for Transmission Control Protocol (TCP) applications (within scope of control of SLA provider)	A, S	% IP address reachable within GTE network at anytime = 97%	✓	✓
Availability	Transport	% IP group (destination address range) reachable for TCP applications within scope of control of SLA provider	A, S	% success of reaching an IP address on extranet of Corp. A between 9:00am-5:00pm = 99%	✓	✓
Availability	Application	Transaction availability	A, S	% success of net.Bank account viewing during 9:00am-11:00pm = 100%	✓	✓
Availability	Application	Unscheduled service Down Time	A, S	% Quake.gameNet downtime = 0	✓	✓
Availability	Transport	Unscheduled network Connections Down Time	A, S, T	Max possible unscheduled RNC-Quest downtime per year = 30 min		✓
Availability	Transport	Initial Connected Speed	A, S	% Initial Connected Speed at 1.1 Mb = 98%		✓
Accessibility	Transport	% rejected calls due to modem problems (e.g., on the PC)	A, S	% Dial-up failure due to modem problem = 7.5%	✓	✓
Accessibility	Transport	% rejected calls due to business policy	A, S	% successful Premium calls over number of committed ports with ports available = 95%	✓	✓
Accessibility	Transport	% rejected calls due to login problems	A, S	Dial-up failure due to login problems = 0%	✓	✓
Accessibility	Transport	% Dial-up attempts/success	A, S	% Successful dial-up attempt = 100%		✓

Class	Sub-class	Performance Metrics	Applicable Probes	Specific Example	Service provider	User
Accessibility	Transport	connection % rejected calls due to IP address unavailability (pool exhaustion)	A, S	% Dial-up failure due to IP address pool exhaustion = 0.001		✓
Accessibility	Transport	% rejected calls due to port unavailability (e.g., lack of modem or network resources)	A, S	% Dial-up failure due to modem ports exhaustion = 0.001		✓
Accessibility	Transport	% rejected calls due to access network unavailability (e.g., network busy, network outage)	A, S	% Dial-up failure due to access network busy = 0.5%		✓
Accessibility	Transport	% rejected calls due to ring no answer	A, S	% Dial-up failure due to ring no answer = 10%		✓
Accessibility	Application	Hosted application accessibility	A, S	% successful net.Phone initiation attempts = 99.9%		✓
Accessibility	Application	Total initiation failures due to application server failure to respond	A, S	% MS Outlook email server sync abnormal abort < 0.002		✓
Accessibility	Application	Total initiation failures due to client abnormally abort	A, S	% MS Outlook email client sync abnormal abort < 0.001		✓
Accessibility	Application	Application destination accuracy (vs. <i>pagelocking</i>)	A, S	% http://ms/complaint hit will lead directly to MS customer service Web page		✓
Reliability	Transport	Mean packet jitters per session	A, S, T	Mean video packet jitters < 2ms	✓	✓
Reliability	Transport	Total discarded packets per session	A, S, T	Max possible packet loss per session < 50	✓	✓
Reliability	Transport	% graceful, % non-graceful disconnections	A, S	% graceful disconnections > 95	✓	✓
Reliability	Application	Mean packet jitters per session	A, S	Mean packet jitter per session < 0.5ms	✓	✓
Reliability	Application	% transaction and sub-transaction errors per transaction type	A, S	% Nortel supplier B inventory query transaction complete successful > 85	✓	✓
Reliability	Application	% transaction and sub-transaction errors per error type	A, S	% Authentication attempt complete successful > 99	✓	✓
Reliability	Application	% transaction and sub-transaction aborted	A, S	% Authentication attempt abort due to server unavailable < 0.2	✓	✓
Reliability	Application	% transaction completion	A, S	% eToys purchase transaction complete successful > 99.9999	✓	✓
Reliability	Transport	Mean packet loss/applications	A, S	Mean video packet loss per Gigabyte transmit or receive < 0.1		✓

Class	Sub-class	Performance Metrics	Applicable Probes	Specific Example	Service provider	User
Reliability	Transport	Mean packet loss/session for subscribers	A, S	Mean net.Meeting packet loss per Gigabyte transmit or receive < 2;		✓
Response Time	Transport	Mean session response time	A, S	Mean net.Meeting packet transmission latency = 90ms		✓
Response Time	Transport	mean session establishment time (inc. AAA, etc.)	A, S	Mean connection to Nortel intranet : 30s		✓
Response Time	Application	Mean transaction end-to-end latency	A, S	Mean email send-to-receive time		✓
Response Time	Application	Mean transaction time per application	A, S	Mean NAT transaction completion time		✓
Response Time	Application	Mean transaction time per type (DNS vs. print spool)	A, S	Mean net.Bank deposit completion time, Mean net.Bank account query completion time		✓
Response Time	Application	Mean sever transaction process time	A, S	Mean Oracle server report response time		✓
Throughput	Transport	Mean bandwidth transmitted over absolute available bandwidth or SLA specified available bandwidth through a given time interval	A, S, T	Mean CIR available = 25 MB; Mean burstable throughput = 5 MB		✓
Throughput	Transport	% VoIP bandwidth of total bandwidth	A, S, T	Mean VoIP CIR available = 7 MB; Mean burstable VoIP throughput = 2 MB		✓
Throughput	Transport	% classes of data bandwidth of total bandwidth	A, S, T	Mean email CIR available = 2 MB; Mean Video CIR = 4 MB		✓
Security	Transport	% encryption (ratio)	A, S	% Packet encryption from subscriber to receiver = 100		✓
Security	Transport	Firewall MTBF	A, S	Firewall MTBF > 99.9%		✓
Security	Transport	Ratio of packets dropped due to security violation (e.g., AH failure, anti-spoofing prevention,...)	A, S	Can we complete?		✓
Security	Transport	Intrusion attempted notification	A, S	Subscribers will be contacted within 30 min. if system detected trouble logins		✓
Security	Transport	% IP blocking correctly	A, S			✓

Class	Sub-class	Performance Metrics	Applicable Probes	Specific Example	Service provider	User
Customer service	Transport	Mean delay from service order to service activation	A, S	Service activation time < 2 days		✓
Customer service	Transport	Mean time to resolution of trouble tickets per class of condition	A, S	Class E1 ticket resolution time = 4 hrs		✓
Customer service	Transport	Mean delay to talking to customer service representative	A, S	Time phone off-hook to customer rep. Contact < 30 sec		✓
Customer service	Transport	Total repeat issues after resolution per year	A, S	Total repeat issues after resolution per year = 0		✓

[0051] The flowchart illustrated at Figure 12 summarizes the operation of the SLA manager 12. At step 200 the

process starts. At step 202 the SLA manager 24 receives performance metrics data from probes T1-n, S1-n and A1-n. At step 204 the SLA manager receives data defining the SLA profile to be followed. At step 206 the SLA manager 24 processes the performance metrics data according to the SLA profile and at step 208 the SLA manager 24 issues one or more SLA actions.

[0052] In a possible variant, the SLA manager can be supplied with simulated performance metrics values generated by a model of the data network, or historical performance metrics data obtained from a real data network. This allows the network operator or service provider to construct hypothetical scenarios and determine future compliance of a data network with current or planned SLA profiles.

[0053] In another variant, the SLA manager 24 does not connect directly to the probes T1-Tn, S1-Sn and A1-An to receive the performance metrics data. Rather the performance metrics data is first recorded on a suitable computer readable storage medium and then played back to the SLA manager 24 that will perform the data processing operations. This procedure is useful in instances where the amount of performance metrics data is so great that real-time processing is not possible or not practical.

[0054] Under another possible variant, the processing service 61, is provided with an information screening capability to deliver only selected information resulting from an SLA action to selected recipients. For example, the service provider or network operator is likely to have more interest in the overall performance of the data network 100, hence he will have access to all the information resulting from the information screening capability to deliver only selected information resulting from an SLA action to selected recipients. For example, the service provider or network operator is likely to have more interest in the overall performance of the data network 100, hence he will have access to all the information resulting from the SLA action(s). In contrast, a certain user 112 - 124 will receive only information reporting if the data network complies with the SLA profile of the user, along with corrective measures taken or planned, if any. The user 112 - 124 is thus prevented from accessing information relative to another user.

[0055] The tool 10 and the SLA manager 24 are implemented in software. The software is run on a suitable computing device that can be standalone or can possess a distributed architecture. An example of such a computing device is shown in Fig. 13. The computing device generally comprises a Central Processing Unit (CPU) 70 that is connected to a memory 72 through a data bus 74. An input/output interface 76, also connected to the data bus 74, allows the computing device to exchange data signals with external devices. The software that implements the functionality of the tool 10 or of the SLA manager 24 is loaded in the memory 72 and executed by the CPU 70. The memory 72 may also include portions reserved for holding data on which the software operates such as the values on which SLA functions or rule apply.

Claims

1. A computer readable storage medium containing a program element suitable for execution by a computing entity for implementing an SLA manager suitable for assessing a performance of a data network, the data network having :
 - a) a transport layer;
 - b) a services layer;
 - c) an applications layer;
 - d) a plurality of probes associated with the transport, services and applications layers of the data network and generating performance metrics data, said SLA manager comprising:
 - i) an input to receive:
 - the performance metrics data;
 - an SLA profile of a user of the data network, the SLA profile defining one or more performance objectives related to a performance of the data network;
 - ii. a processing service for measuring a certain performance objective defined by the SLA profile at least in part on the basis of performance metrics data from a first selected one of the plurality of probes associated with one layer of the data network in conjunction with performance metrics data from a second selected one of the plurality of probes associated with another layer of the data network;
 - iii. said processing service being operative to generate at least one SLA action in response to the measuring.
2. A computer readable storage medium as defined in claim 1, wherein said processing service is operative for measuring a certain performance objective defined by the SLA profile at least in part on the basis of performance metrics data from a probe associated with the transport layer of the data network in conjunction with performance metrics

data from a probe associated with the services layer of the data network.

3. A computer readable storage medium as defined in claim 1, wherein said processing service is operative for measuring a certain performance objective defined by the SLA profile at least in part on the basis of performance metrics data from a probe associated with the applications layer of the data network in conjunction with performance metrics data from a probe associated with the services layer of the data network.
4. A computer readable storage medium as defined in claim 1, wherein said processing service is operative for measuring a certain performance objective defined by the SLA profile at least in part on the basis of performance metrics data from a probe associated with the transport layer of the data network in conjunction with performance metrics data from a probe associated with the services layer of the data network in conjunction with performance metrics data from a probe associated with the applications layer of the data network.
5. A computer readable storage medium as defined in claim 1, wherein the performance metrics data is real time data.
6. A computer readable storage medium as defined in claim 1, wherein the performance metrics data is historical data.
7. A computer readable storage medium as defined in claim 2, wherein the certain performance objective is selected from the group consisting of network element MTBF, network segment MTBF and network element MTTR, network segment MTTR.
8. A computer readable storage medium as defined in claim 3, wherein the certain performance objective is selected from the group consisting of % of IP destination address reachable for a TCP application, % of IP group reachable for a TCP application, transaction availability, unscheduled service down time, initial connection speed, % of rejected calls due to modem problem, % of rejected calls due to business policy, % of rejected calls due to login problems, % of dial-up attempts/success connection, % of rejected calls due to IP address unavailability, % of rejected calls due to port unavailability, % of rejected calls due to access network unavailability, % of rejected calls due to ring no answer, hosted application accessibility, total initiation failures due to application server failure to respond, total initiation failures due to client abnormally abort, application destination accuracy, % of graceful disconnections, % of non-graceful disconnections, mean packet jitters per session, % transaction and sub-transaction errors per transaction type, % transaction and sub-transaction errors per error type, % transactions and sub-transactions aborted, % transaction completion, mean packet loss per application, mean packet loss per session per subscriber, mean session response time, mean session establishment time, mean transaction end-to-end latency, mean transaction time per application, mean transaction time per type, mean server transaction process time, % of encryption ratio, firewall MTBF, ratio of packets dropped due to security violation, intrusion attempted notification, % IP blocking correctly, mean delay from service order to service activation, mean time to resolution of trouble tickets per class of condition and mean delay to talking to customer service representative and total repeat issues after resolution per year.
9. A computer readable storage medium as defined in claim 4, wherein the certain performance objective is selected from the group consisting of unscheduled network connections down time, mean packet jitters per session, total number of packets discarded per session, mean bandwidth transmitted over absolute available bandwidth, SLA specified available bandwidth through a given time interval, % of VoIP bandwidth of total bandwidth and % classes of data bandwidth of total bandwidth.
10. A computer readable storage medium as defined in claim 1, wherein the SLA action is selected in the group consisting of positive SLA action and null SLA action.
11. A computer readable storage medium as defined in claim 10, wherein said positive SLA action is selected in the group consisting of creation of a report, sending a command to an e-mail server and creating, modifying or deleting a billing record.
12. An SLA manager suitable for assessing a performance of a data network, the data network having :
 - a) a transport layer;
 - b) a services layer;
 - c) an applications layer;
 - d) a plurality of probes associated with the transport, services and applications layers of the data network and

generating performance metrics data, said SLA manager comprising:

i) an input to receive:

- the performance metrics data;
- an SLA profile of a user of the data network, the SLA profile defining one or more performance objectives related to a performance of the data network;

ii) a processing service for measuring a certain performance objective defined by the SLA profile at least in part on the basis of performance metrics data from a first selected one of the plurality of probes associated with one layer of the data network in conjunction with performance metrics data from a second selected one of the plurality of probes associated with another layer of the data network; and

iii) said processing service being operative to generate at least one SLA action in response to the measuring.

13. A method to assess a performance of a data network, including:

- a) a transport layer;
 - b) a services layer;
 - c) an applications layer;
 - d) a plurality of probes associated with the transport, services and applications layers of the data network and generating performance metrics data;
- said method comprising:

i. measuring a certain performance objective defined by a predetermined SLA profile at least in part on the basis of performance metrics data from a first selected one of the plurality of probes associated with one layer of the data network in conjunction with performance metrics data from a second selected one of the plurality of probes associated with another layer of the data network;

ii. generating at least one SLA action in response to the measuring.

14. A computer readable storage medium containing a program element suitable for execution by a computing device, said program element implementing a tool to create an SLA profile for processing by an SLA manager to assess a performance of a data network, the data network having:

- a) a transport layer;
- b) a services layer;
- c) an applications layer;
- d) a plurality of probes associated with the transport, services and applications layers of the data network and generating performance metrics data;
- e) said tool comprising:

a first unit allowing an operator to specify through an operator interface at least one SLA performance metric measurable by assessing performance metrics data at least in part on the basis of performance metrics data from a first selected one of the plurality of probes associated with one layer of the data network in conjunction with performance metrics data from a second selected one of a plurality of probes associated with another layer of the data network;

i. a second unit allowing the user to specify through the user interface at least one SLA rule;

ii. said program element being operative to generate in response to the specification of the at least one SLA performance metric and the specification of the at least one SLA rule by the operator the SLA profile; and

iii. an output to release the SLA profile.




FIG. 1




FIG. 2

FIG. 3B





FIG. 3A




FIG. 4




FIG. 5




FIG. 6

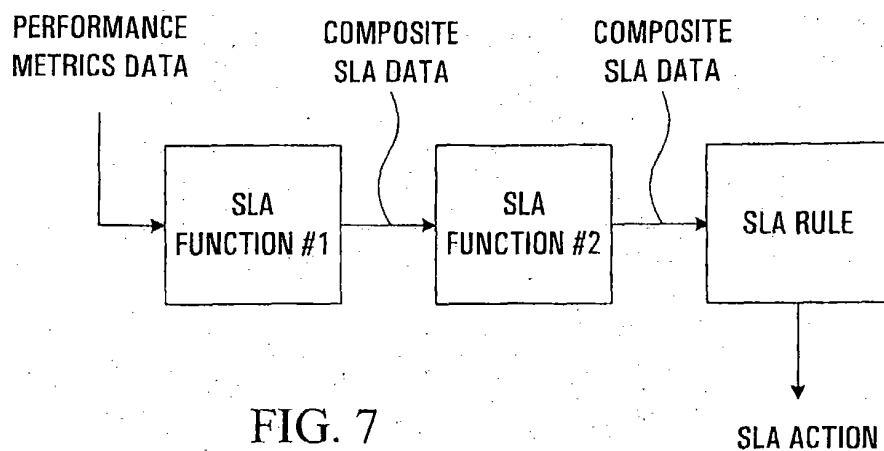


FIG. 7




FIG. 8




FIG. 9A




FIG. 9B




FIG. 10




FIG. 11

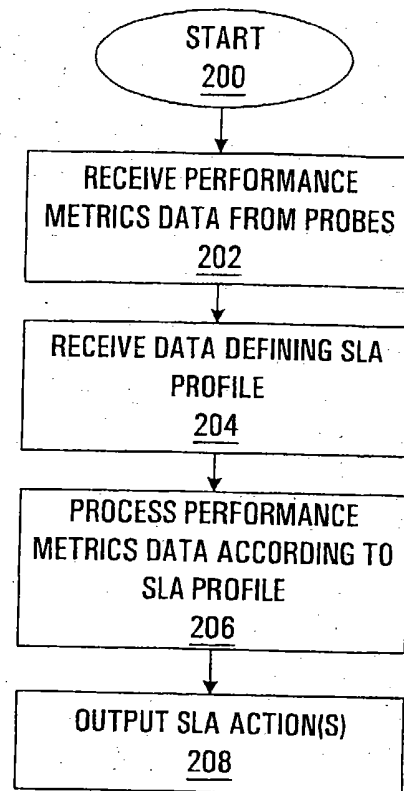


FIG. 12




FIG. 13